Quantisation Noise in Digital Camera Images

Dr. Friedrich Dierks, Basler AG

Abstract *This paper analyzes the effect of quantisation noise on the accuracy of measurements in digital camera images.*

Digital cameras contain an imaging sensor which delivers an analogue signal. This signal is converted by an analog-to-digital converter (ADC) into a stream of digital numbers with q bits resolution. If the noise on the analogue signal is in the same order of magnitude or smaller than the quantisation step Δy of the ADC, the quantisation procedure creates errors which can be described by quantisation noise added to the signal. This paper deals with the question how to choose the quantisation step so that the quantisation noise is negligible compared to the noise contained in the analogue signal.

Standard Model for Quantisation Noise

The quantisation process is highly non-linear and thus difficult to describe analytically. Different approaches have been reported in literature. The most widely used approximation is based on the assumption that the signal is stochastic and that all amplitude values will be found with the same probability in the signal (see for example [**2**]). As a result the quantisation noise can be described by the mean

$$\mu_q = \mu \tag{1}$$

a the variance

$$\sigma_q^2 = \sigma^2 + \frac{1}{12} \tag{2}$$

In these formulas μ is the mean and σ the standard deviation of the analogue signal. μ_q and σ_q are the respective values of the quantised signal. Note that since the analogue values are given in quantised units the quantisation step is by definition $\Delta y = 1$.

As shown in [**2**] equations (1) and (2) are true even for a more general class of signals. The quantisation of the amplitude of a signal can be described as sampling process. As an example Figure 1 shows a constant signal with added Gaussian noise. Sampling the signal values results in an impulse train where the impulses have a distance of Δy.

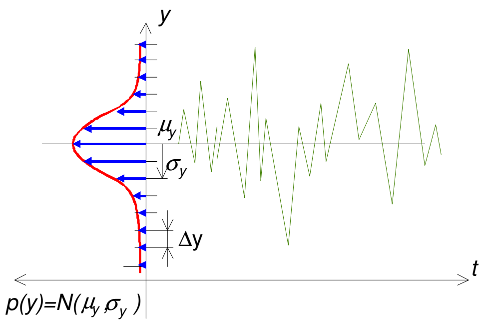

Figure 1 Probability of sampled amplitude values

From the sampling theorem follows: If the Fourier transformation of the probability function is bandwidth limited, equations (1) and (2) hold true exactly.

Unfortunately the most widely used signal model for noisy images, a virtually constant signal with added Gaussian white noise, is not bandwidth limited. So for the signal depicted in Figure 1 equations (1) and (2) cannot be applied.

Probability Function for Samples Gaussian Noise

In order to find a better description of the quantisation effects first the probability distribution of the sampled signal is computed. The probability that the analogue signal's amplitude is located in the range [y, y+dy] can be written as

$$dp = \varphi(y;\mu;\sigma)dy \tag{3}$$

with

$$\varphi(y,\mu,\sigma) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{y-\mu}{\sigma}\right)^2} \tag{4}$$

where μ is the mean and σ is the standard deviation of the signal. Integrating this equation results in the probability that the signal amplitude is located in the range [$-\infty$, y]:

$$p_{y\in[-\infty,y]} = \Phi(y,\mu,\sigma) = \int_{-\infty}^{y} \varphi(\xi,\mu,\sigma)d\xi \qquad (5)$$

This integral cannot be solved analytically but only numerically.

A convenient way to get the numerical values for the probability distribution and it's integral are the following functions provided by Excel:

$\varphi(y,\mu,\sigma)$ = `NORMVERT(y;mean;sigma;FALSCH)`
$\Phi(y,\mu,\sigma)$ = `NORMVERT(y;mean;sigma;WAHR)`

The quantisation can be described by the characteristic curve shown in Figure 2.

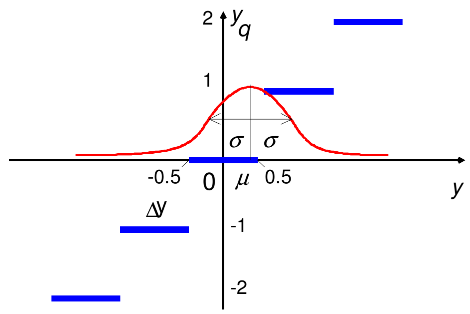

Figure 2 Probability of sampled amplitude values

The probability distribution of the quantised signal y_q is:

$$\begin{aligned} p_{y_q=n} &= p_q(n) \\ &= \Phi(n+0.5,\mu,\sigma) - \Phi(n-0.5,\mu,\sigma) \end{aligned} \qquad (6)$$

The mean of the quantised signal can be computed as

$$\mu_q = \sum_{i=-\infty}^{+\infty} i \cdot p_q(i) \qquad (7)$$

and the variance as

$$\sigma_q^2 = \sum_{i=-\infty}^{+\infty} (i-\mu_q)^2 p_q(i) \qquad (8)$$

To simplify the analysis two restrictions are introduced:

- For symmetry reasons it is sufficient to deal with the range $\mu \in [-0.5,+0.5]$
- As we will see quantisation effects for $\sigma > 2$ can be neglected so we deal with $\sigma \leq 2$ only.

Under these conditions we can truncate the sums and approximate:

$$\mu_q \approx \sum_{i=-5}^{+5} i \cdot p_q(i) \qquad (9)$$

and the variance as

$$\sigma_q^2 \approx \sum_{i=-5}^{+5} (i-\mu_q)^2 p_q(i) \qquad (10)$$

These two formulas can be – with some typing – easily implemented in Excel. The are the base of the analysis in the rest of this article.

How the Mean is Affected by Quantisation

Figure 3 shows how quantisation affects the mean of the quantised signal: The mean of the analogue signal varies in the range $-0.5 \leq \mu \leq 0.5$; the standard deviation is constant $\sigma = 0.1$. The black line labeled *m* shows the mean μ of the analogue signal; the blue curve labeled *M* shows the mean $\mu_{y.q}$ of the quantised signal; the red line labeled *M-m* shows the error. Note that the error is substantial and does not follow the prediction of the standard error model for amplitude bandwidth limited signals which predicts that the error is always zero.

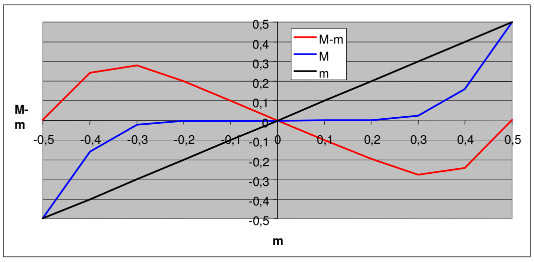

Figure 3 Quantisation error of the mean (for details see text)

To understand the error mechanism look at the three cases A-C drawn in Figure 4. In case A the noise is centered around the switching point from $y_q = -1$ to $y_q = 0$. As a consequence the mean of the output value will be 0.5 which is also the mean of the analogue signal. This corresponds with the observance that the error is always zero at $\mu = \pm 0.5$.

In case B the noise is in the middle of the step. The output values will centered around zero. This corresponds with the observance that the error is always zero at $\mu = 0$.

In case C the noise is placed between $\mu = 0$ and $\mu = 0.5$. If the standard deviation is small the output signal will be mainly zero and thus deviate from $\mu > 0$. For larger standard deviations the error vanishes approximately.

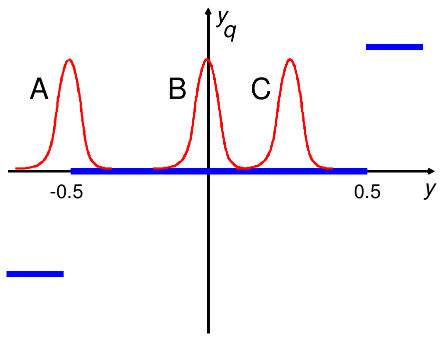

Figure 4 Error mechanism

How the Standard Deviation is Affected by Quantisation

Figure 5 shows how quantisation affects the standard deviation of the quantised signal: The mean is constant $\mu = -0.49$; the standard deviation of the analogue signal varies in the range $0 \le \sigma \le 1$;. The black line labeled *s* shows the standard deviation σ of the analogue signal; the blue curve labeled *S* shows the standard deviation σ_q of the quantised signal; the red line labeled *S-s* shows the error; the green line shows the error predicted by the standard formula for amplitude bandwidth limited noise.

$$\sigma_q^2 = \sigma^2 + \frac{1}{12} \tag{11}$$

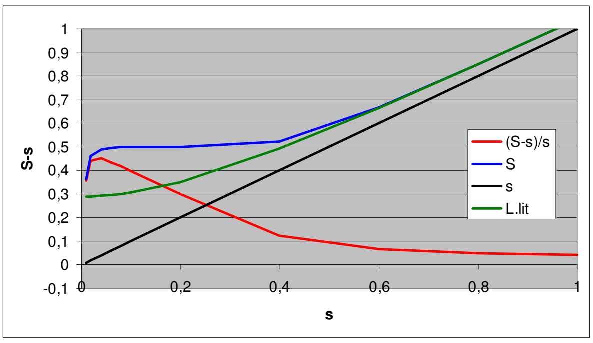

Figure 5 Quantisation error of the standard deviation and $\mu = -0.49$ (for details see text)

With a mean of $\mu = -0.49$ the situation is similar to case A in Figure 4. Even if the standard deviation of the analogue noise is small the output signal will toggle between $y_q = -1$ to $y_q = 0$. As a result the standard deviation of the quantised signal saturates to $\sigma_q = 0.5$ for $\sigma < 0.4$. This is in contrast to the standard model which will saturate more slowly to $\sigma_q = 0.3$.

Figure 6 shows the error of the standard deviation for case B with $\mu = 0$. For small standard deviations of the analogue signal the noise of the digitized signal drops to zero. This is plausible because the output signal will stay mainly constant at $y_q = 0$. Note that this behavior is also in contradiction to the predictions of the standard model.

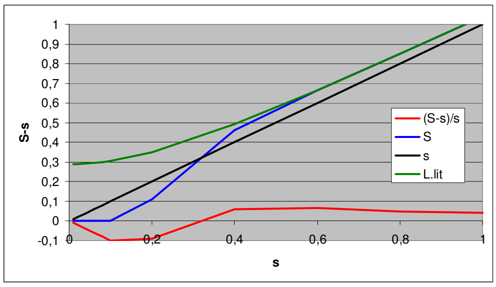

Figure 6 Quantisation error of the standard deviation and $\mu = 0$ (for details see text)

Choosing the Quantisation Step Size

The quantisation step Δy should be chosen small enough that the quantisation error can be neglected. In Figure 7 the blue line shows the maximum error of the mean versus the noise level of the analogue signal; the red line shows the maximum error of the standard variation. The maximum was in both cases taken with respect to the mean of the analogue signal which was varied in the range $-0.5 \le \mu \le 0.5$. The green line shows the error of the standard deviation predicted by the standard model.

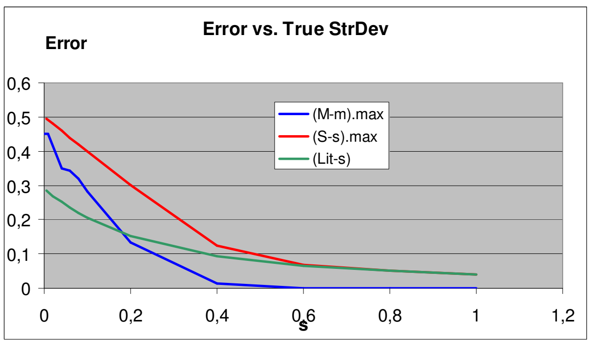

Figure 7 Quantisation error versus standard deviation of the analogue signal (for details see text)

From the picture follows as a general rule that the quantisation step should be chosen as

$$\Delta y \leq 2\sigma_y \tag{12}$$

where σ_y is the variance of the analogue signal. If this condition is fulfilled hardly any error in the mean will take place. In addition the standard deviation will follow the standard formula

$$\sigma_q^2 = \sigma^2 + \frac{1}{12} \tag{13}$$

and the quantisation noise will be <17%.

For smaller quantisation steps there will be a strong error in both the mean and the standard deviation which heavily depends on the mean of the analogue signal and is not predictable by a linear model.

An interesting corollary is that an analogue signal with p bits dynamic range can be transferred using p-1 digital bits.

Literature

[**1**] Dierks, F.; “Comparing Digital Cameras with Respect to Sensitivity and Noise”; Basler internal paper

[**2**] Kronmüller, Heinz; Digitale Signalverarbeitung, Springer 1991, ISBN 3-540-54128-4